Remarks

Entrance of this amendment and allowance of the pending claims are respectfully requested. Claims 12-20 remain pending.

By this paper, the system, program storage device and data structure claim sets are canceled without prejudice to the re-filing thereof in one or more continuation or divisional applications. Applicants are not conceding in this application that these canceled claims are not patentable over the art cited in the Office Action, but rather are submitting the claim cancellations to place all method claims in one patent and move the other classes of statutory subject matter to one or more continuation or divisional patent applications. Applicants respectfully reserve the right to pursue these canceled claims in one or more continuation or divisional patent applications.

In pending claims 12-20, independent claim 12 is amended to more particularly point out, and distinctly claim certain aspects of the present invention. Specifically, this claim is amended to recite a processing method for a distributed *parallel computing system.* The processing method utilizes a dedicated collective offload engine, *which is a hardware device* coupled to the switch fabric. *The hardware device is a specialized device dedicated to providing collective processing in hardware of data from the processing nodes* of the distributed, parallel computing system. Support for the amended claims can be found throughout the application as filed. For example, reference specification paragraphs [0014] & [0015]. No new matter is added to the application by any amendment presented.

Claims 1-6, 9-17, 20-26, 29-30 were initially rejected under 35 U.S.C. § 102(e) as being anticipated by Burianek et al. (U.S. Patent No. 7,082,457; hereinafter Burianek), while claims 7, 8, 18, 19, 27 & 28 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Burianek. These rejections are respectfully traversed to any extent deemed applicable to the claims presented herewith, and reconsideration thereof is requested for the reasons set forth below.

As amended, Applicants recite *a data processing method implemented within a distributed parallel computing system.* Burianek does not describe a distributed parallel computing system, but rather describes a client server computing environment. In view of this

basic difference, Applicants respectfully submit that the data processing method recited in their pending claims, is not anticipated by the teachings of Burianek.

In addition, Applicants recite providing, by a *dedicated collective offload engine* coupled to a switch fabric in a distributed parallel computing system, collective processing of data. There is no dedicated collective offload engine in Burianek as the term is employed in Applicants' specification and claims. In the Office Action, Applicants recited dedicated collective offload engine is analogized to server 215 of Burianek. This analogy is respectfully traversed.

Server 215 in Burianek is described as a project management central server that directs signals sent to and from the components of the distributed computing environment. This server includes a delegation component which sends and retrieves information about project tasks stored in the database 210. Thus, server 215 in Burianek is a conventional server system. This is distinguished from Applicants' *dedicated collective offload engine* which provides collective processing of data. In Applicants' invention, the *dedicated collective offload engine is a hardware device* coupled to the switch fabric. This hardware device (previously recited in dependent claim 13), distinguishes Applicants' invention from Burianek. In Burianek, the processing described is implemented in software. In contrast, Applicants' processing is implemented in hardware, that is, in the hardware device which is the *dedicated collective offload engine* (one embodiment of which is depicted in FIG. 2 of the present application).

With respect to the subject matter of original claim 13, and in particular, the *dedicated collective offload engine* being implemented *as a hardware device*, the Office Action references the remote computer of FIG. 1 in Burianek. The Office Action asserts that the remote computer is a hardware device. This characterization of the remote computer is respectfully traversed. As is well known in the art, a computer comprises both hardware and software. There is no discussion in Burianek that the remote computer 11 of FIG. 1 is a hardware device only. This difference between Burianek and Applicants' recited invention is further characterized in independent claim 12 presented herewith, wherein it is recited that *the hardware device is a specialized device dedicated to providing collective processing in hardware of data from the at least some processing nodes*. In Applicants' invention, the collective processing is implemented in hardware within the specialized hardware device coupled to the switch fabric. This

specialized hardware device is referred to as the dedicated collective offload engine. As recited in Applicants' independent claim 12, the specialized device is dedicated to providing collective processing in hardware of the data from the at least some processing nodes. There is no such specialized device described in Burianek. For at least this additional reason, Applicants respectfully submit that independent claim 12 patentably distinguishes over the applied and known art.

Still further, Applicants recite a data processing method which includes *collective processing* of data from the at least some processing nodes of the multiple processing nodes of the distributed, parallel computing system. In amended claim 12, collective processing is recited to implement a *collective operation* on the data from the at least some processing nodes. The phrases *collective processing* and *collective operation* are terms of art which refer to a particular type of data processing. A collective operation is conventionally an arithmetic operation executed across data from multiple nodes of a distributed, parallel computing system.

As explained in Applicants' Background of the Invention, implementation of collective processing typically includes using a software tree approach, wherein message passing facilities are used to form a virtual tree of processes. A drawback to this approach is the serialization of delays at each stage of the tree. These delays are additive in the overall overhead associated with the collective processing. Furthermore, this software tree approach results in a theoretical logarithmic scaling latency of the overall collective processing versus system size. Due to interference from daemons, interrupts and other background activity, cross traffic, and the unsynchronized nature of independent operating system images and their dispatch cycles, measured values of scaling latency are usually significantly worse than theoretical values. Responsive to this issue, Applicants describe a novel *collective processing approach* which mitigates the large latency associated with the software tree implementation. In Applicants' approach, a dedicated collective offload engine (which is a hardware device, coupled to the switch fabric) is employed to provide the collective processing of data from the multiple processing nodes. Applicants' hardware device is a specialized device dedicated to providing the collective processing in hardware of the data, *and the collective processing implements a collective operation on the data.* (As recited in dependent claim 20, this advantageously avoids the need for a software tree.)

An internet search on the phrase "collective operation" in a distributed parallel computing system, or "collective processing" provides support for the above-noted meaning of these phrases as employed in the art. Applicants respectfully request that this meaning be given consideration when evaluating the claims at issue. Burianek does not describe collective processing *per se*, nor is a collective operation as the term is understood in the art, described in Burianek. As such, independent claim 12 patentably distinguishes over the applied art.

For at least the above-noted reasons, Applicants respectfully request reconsideration and withdrawal of the rejection to independent claim 12 presented herewith.

The dependent claims are believed allowable for the same reasons as the independent claims, as well as for their own additional characterizations.

For example, amended claim 13 recites that the collective operation is a Message Passing Interface (MPI) collective operation. An MPI collective operation is a particular type of collective operation implemented within the MPI standard. Details on MPI collective operations are provided at http://www.redbooks.ibm.com/redbooks/pdfs/sg245380.pdf. For example, reference chapter 2 thereof. There is no discussion in Burianek of the MPI standard, or of a collective operation implemented within the standard. It is Applicants' collective processing employing the dedicated collective offload engine (i.e., specialized hardware device) which allows for collective processing in hardware of the data from the multiple processing nodes of the distributed, parallel computing system. No such device is taught or suggested in the art of record.

Claim 20 specifies that the collective processing of Applicants' data processing system executes the collective operation for the at least some processing nodes *without using a software tree*. As noted above, a software tree is conventionally employed to implement collective processing within a distributed, parallel computing system. Applicants' recited collective processing in hardware accomplishes execution of the collective operation without using a software tree. The art of record does not describe such a protocol.

All claims are believed to be in condition for allowance, and such action is respectfully requested.

Should the Examiner have any reservation regarding the patentability of the claims presented, however, Applicants' undersigned representative respectfully requests the opportunity for an Examiner Interview to discuss the claims in the hope of advancing prosecution of this application.

Respectfully submitted,

Kevin P. Radigan, Esq.
Attorney for Applicants
Registration No.: 31,789

Dated: May 9, 2008.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579